WHAT IS CLAIMED IS:

1. A computer-implemented method of managing a storage environment comprising a plurality of storage units, the method comprising:

detecting a condition associated with a first storage unit from the plurality of storage units;

determining a first group from a plurality of groups to which the first storage unit belongs, wherein each group comprises one or more storage units from the plurality of storage units and inclusion of a storage unit in a group depends on a cost of storing data on the storage unit;

identifying a second group from the plurality of groups having an associated data storage cost that is lower than a data storage cost associated with the first group;

identifying a file stored on the first storage unit to be moved;

identifying a storage unit from the second group for storing the file; and

moving the file from the first storage unit to the storage unit from the second group that has been identified for storing the file.

2. The method of claim 1 further comprising repeating, the identifying a file stored on the first storage unit to be moved, the identifying a storage unit from the second group for storing the file, and the moving the file from the first storage unit to the storage unit from the second group that has been identified for storing the file, until the condition is resolved.

3. The method of claim 2 wherein the first storage unit stores a set of migrated files and a set of original files, the set of migrated files comprising files that have been migrated or remigrated from their original storage locations, the set of original files comprising files that have not been migrated from their original storage locations, and wherein a file from the set of original files is not selected to be moved until all files in the set of migrated files have been selected and moved from the first storage unit.

4. The method of claim 2 wherein detecting a condition associated with the first storage unit comprises detecting that used storage capacity for the first storage unit has exceeded a first threshold, and the condition is resolved when the used storage capacity for the first storage unit does not exceed the first threshold.

5. The method of claim 1 wherein identifying a storage unit from the second group comprises identifying a storage unit from one or more storage units in the second group that is least full.

6. The method of claim 1 wherein identifying a storage unit from the second group comprises:

generating a score for each storage unit in the second group; and

selecting a storage unit from the second group based upon the scores generated for the one or more storage units in the second group.

7. The method of claim 1 wherein the first storage unit stores a plurality of files and identifying a file stored on the first storage unit to be moved comprises:

generating a score for each file in the plurality of files stored on the first storage unit; and

selecting a file to be moved from the plurality of files based upon the scores generated for the files in the plurality of files.

8. The method of claim 1 wherein the first storage unit is assigned to a first server and the storage unit from the second group to which the file from the first storage unit is moved is assigned to a second server distinct from the first server.

9. A computer-implemented method of managing a storage environment comprising a plurality of storage units, the method comprising:

detecting a condition associated with a first storage unit from the plurality of storage units;

identifying a file stored on the first storage unit to be moved;

identifying a storage unit from the plurality of storage units for storing the file, wherein the data storage cost associated with identified storage unit is lower than a data storage cost associated with the first storage unit; and

moving the file from the first storage unit to the storage unit from the second group that has been identified for storing the file.

10. The method of claim 9 wherein identifying a storage unit from the plurality of storage units for storing the file comprises:

identifying a set of storage units from the plurality of storage units that have an associated data storage cost that is lower than the data storage cost associated with the first storage unit; and

selecting a storage unit for storing the file from the set of storage units.

11. The method of claim 9 further comprising repeating, the identifying a file stored on the first storage unit to be moved, the identifying a storage unit from the plurality of storage units for storing the file, and the moving the file from the first storage unit to the storage unit from the second group that has been identified for storing the file, until the condition is resolved.

12. The method of claim 11 wherein detecting a condition associated with the first storage unit comprises detecting that used storage capacity for the first storage unit has exceeded a first threshold, and the condition is considered resolved when the used storage capacity for the first storage unit does not exceed the first threshold.

13. The method of claim 11 wherein the first storage unit stores a set of migrated files and a set of original files, the set of migrated files comprising files that have been migrated or remigrated from their original storage locations, the set of original files comprising files that have not been migrated from their original storage locations, and wherein a file from the set of original files is not selected to be moved until all files in the set of migrated files have been selected and moved from the first storage unit.

14. The method of claim 9 wherein identifying a storage unit from the plurality of storage units for storing the file comprises identifying a storage unit from the plurality of storage units that is least full.

15. The method of claim 9 wherein identifying a storage unit from the plurality of storage units for storing the file comprises:

generating scores for storage units in the plurality of storage units; and

selecting a storage unit from the plurality of storage units based upon the generated scores.

16. The method of claim 9 wherein the first storage unit stores a plurality of files and identifying a file stored on the first storage unit to be moved comprises:

generating a score for each file in the plurality of files stored on the first storage unit; and

selecting a file to be moved from the plurality of files based upon the scores generated for the files in the plurality of files.

17. The method of claim 9 wherein the first storage unit is assigned to a first server and the storage unit from the plurality of storage units to which the file from the first storage unit is moved is assigned to a second server distinct from the first server.

18. A computer program product stored on a computer-readable storage medium for managing a storage environment comprising a plurality of storage units, the computer program product comprising:

code for detecting a condition associated with a first storage unit from the plurality of storage units;

code for determining a first group from a plurality of groups to which the first storage unit belongs, wherein each group comprises one or more storage units from the plurality of storage units and inclusion of a storage unit in a group depends on a cost of storing data on the storage unit;

code for identifying a second group from the plurality of groups having an associated data storage cost that is lower than a data storage cost associated with the first group;

code for identifying a file stored on the first storage unit to be moved;

code for identifying a storage unit from the second group for storing the file; and

code for moving the file from the first storage unit to the storage unit from the second group that has been identified for storing the file.

19. The computer program product of claim 18 further comprising code for repeating, the identifying a file stored on the first storage unit to be moved, the identifying a storage unit from the second group for storing the file, and the moving the file from the first storage unit to the storage unit from the second group that has been identified for storing the file, until the condition is resolved.

20. The computer program product of claim 19 wherein the first storage unit stores a set of migrated files and a set of original files, the set of migrated files

comprising files that have been migrated or remigrated from their original storage locations, the set of original files comprising files that have not been migrated from their original storage locations, and wherein a file from the set of original files is not selected to be moved until all files in the set of migrated files have been selected and moved from the first storage unit.

21. The computer program product of claim 19 wherein the code for detecting a condition associated with the first storage unit comprises code for detecting that used storage capacity for the first storage unit has exceeded a first threshold, and the condition is resolved when the used storage capacity for the first storage unit does not exceed the first threshold.

22. The computer program product of claim 18 wherein the code for identifying a storage unit from the second group comprises code for identifying a storage unit from one or more storage units in the second group that is least full.

23. The computer program product of claim 18 wherein the code for identifying a storage unit from the second group comprises:

code for generating a score for each storage unit in the second group; and

code for selecting a storage unit from the second group based upon the scores generated for the one or more storage units in the second group.

24. The computer program product of claim 18 wherein the first storage unit stores a plurality of files and the code for identifying a file stored on the first storage unit to be moved comprises:

code for generating a score for each file in the plurality of files stored on the first storage unit; and

code for selecting a file to be moved from the plurality of files based upon the scores generated for the files in the plurality of files.

25. The computer program product of claim 18 wherein the first storage unit is assigned to a first server and the storage unit from the second group to which the file from the first storage unit is moved is assigned to a second server distinct from the first server.

26. A computer program product stored on a computer-readable storage medium for managing a storage environment comprising a plurality of storage units, the computer program product comprising:

code for detecting a condition associated with a first storage unit from the plurality of storage units;

code for identifying a file stored on the first storage unit to be moved;

code for identifying a storage unit from the plurality of storage units for storing the file, wherein the data storage cost associated with identified storage unit is lower than a data storage cost associated with the first storage unit; and

code for moving the file from the first storage unit to the storage unit from the second group that has been identified for storing the file.

27. The computer program product of claim 26 wherein the code for identifying a storage unit from the plurality of storage units for storing the file comprises:

code for identifying a set of storage units from the plurality of storage units that have an associated data storage cost that is lower than the data storage cost associated with the first storage unit; and

code for selecting a storage unit for storing the file from the set of storage units.

28. The computer program product of claim 26 further comprising code for repeating, the identifying a file stored on the first storage unit to be moved, the identifying a storage unit from the plurality of storage units for storing the file, and the moving the file from the first storage unit to the storage unit from the second group that has been identified for storing the file, until the condition is resolved.

29. The computer program product of claim 28 wherein the code for detecting a condition associated with the first storage unit comprises code for detecting that used storage capacity for the first storage unit has exceeded a first threshold, and the condition is considered resolved when the used storage capacity for the first storage unit does not exceed the first threshold.

30. The computer program product of claim 28 wherein the first storage unit stores a set of migrated files and a set of original files, the set of migrated files comprising files that have been migrated or remigrated from their original storage locations,

the set of original files comprising files that have not been migrated from their original storage locations, and wherein a file from the set of original files is not selected to be moved until all files in the set of migrated files have been selected and moved from the first storage unit.

31. The computer program product of claim 26 wherein the code for identifying a storage unit from the plurality of storage units for storing the file comprises code for identifying a storage unit from the plurality of storage units that is least full.

32. The computer program product of claim 26 wherein the code for identifying a storage unit from the plurality of storage units for storing the file comprises:

code for generating scores for storage units in the plurality of storage units; and

code for selecting a storage unit from the plurality of storage units based upon the generated scores.

33. The computer program product of claim 26 wherein the first storage unit stores a plurality of files and the code for identifying a file stored on the first storage unit to be moved comprises:

code for generating a score for each file in the plurality of files stored on the first storage unit; and

code for selecting a file to be moved from the plurality of files based upon the scores generated for the files in the plurality of files.

34. The computer program product of claim 26 wherein the first storage unit is assigned to a first server and the storage unit from the plurality of storage units to which the file from the first storage unit is moved is assigned to a second server distinct from the first server.

35. A system comprising:

a plurality of storage units; and

a data processing system configured to manage the plurality of storage units, wherein the data processing system is configured to:

detect a condition associated with a first storage unit from the plurality of storage units;

determine a first group from a plurality of groups to which the first storage unit belongs, wherein each group comprises one or more storage units from the plurality of storage units and inclusion of a storage unit in a group depends on a cost of storing data on the storage unit;

identify a second group from the plurality of groups having an associated data storage cost that is lower than a data storage cost associated with the first group;

identify a file stored on the first storage unit to be moved;

identify a storage unit from the second group for storing the file; and

move the file from the first storage unit to the storage unit from the second group that has been identified for storing the file.

36. The system of claim 35 wherein the data processing system is configured to repeat, the identification of a file stored on the first storage unit to be moved, the identification of a storage unit from the second group for storing the file, and the move of the file from the first storage unit to the storage unit from the second group that has been identified for storing the file, until the condition is resolved.

37. The system of claim 36 wherein the first storage unit stores a set of migrated files and a set of original files, the set of migrated files comprising files that have been migrated or remigrated from their original storage locations, the set of original files comprising files that have not been migrated from their original storage locations, and wherein a file from the set of original files is not selected to be moved until all files in the set of migrated files have been selected and moved from the first storage unit.

38. The system of claim 36 wherein the data processing system is configured to detect that used storage capacity for the first storage unit has exceeded a first threshold, and the condition is resolved when the used storage capacity for the first storage unit does not exceed the first threshold.

39. The system of claim 35 wherein the data processing system is configured to identify a storage unit from one or more storage units in the second group that is least full as the storage unit for storing the file.

40. The system of claim 35 wherein the data processing system is configured to:

generate a score for each storage unit in the second group; and

select a storage unit from the second group based upon the scores generated for the one or more storage units in the second group.

41. The system of claim 35 wherein the first storage unit stores a plurality of files and the data processing system is configured to:

generate a score for each file in the plurality of files stored on the first storage unit; and

select a file to be moved from the plurality of files based upon the scores generated for the files in the plurality of files.

42. The system of claim 35 wherein the first storage unit is assigned to a first server and the storage unit from the second group to which the file from the first storage unit is moved is assigned to a second server distinct from the first server.

43. A system comprising:

a plurality of storage units; and

a data processing system configured to manage the plurality of storage units, wherein the data processing system is configured to:

detect a condition associated with a first storage unit from the plurality of storage units;

identify a file stored on the first storage unit to be moved;

identify a storage unit from the plurality of storage units for storing the file, wherein the data storage cost associated with identified storage unit is lower than a data storage cost associated with the first storage unit; and

move the file from the first storage unit to the storage unit from the second group that has been identified for storing the file.

44. The system of claim 43 wherein the data processing system is configured to:

identify a set of storage units from the plurality of storage units that have an associated data storage cost that is lower than the data storage cost associated with the first storage unit; and

select a storage unit from the set of storage units for storing the file.

45. The system of claim 43 wherein the data processing system is configured to repeat, the identification of a file stored on the first storage unit to be moved, the identification of a storage unit from the plurality of storage units for storing the file, and the move of the file from the first storage unit to the storage unit from the second group that has been identified for storing the file, until the condition is resolved.

46. The system of claim 45 wherein the data processing system is configured to detect that used storage capacity for the first storage unit has exceeded a first threshold, and the condition is considered resolved when the used storage capacity for the first storage unit does not exceed the first threshold.

47. The system of claim 45 wherein the first storage unit stores a set of migrated files and a set of original files, the set of migrated files comprising files that have been migrated or remigrated from their original storage locations, the set of original files comprising files that have not been migrated from their original storage locations, and wherein a file from the set of original files is not selected to be moved until all files in the set of migrated files have been selected and moved from the first storage unit.

48. The system of claim 43 wherein the data processing system is configured to identify a storage unit from the plurality of storage units that is least full as the storage unit for storing the file.

49. The system of claim 43 wherein the data processing system is configured to:

generate scores for storage units in the plurality of storage units; and

select a storage unit from the plurality of storage units based upon the generated scores.

50. The system of claim 43 wherein the first storage unit stores a plurality of files and the data processing system is configured to:

generate a score for each file in the plurality of files stored on the first storage unit; and

select a file to be moved from the plurality of files based upon the scores generated for the files in the plurality of files.

51. The system of claim 43 wherein the first storage unit is assigned to a first server and the storage unit from the plurality of storage units to which the file from the first storage unit is moved is assigned to a second server distinct from the first server.

52. A system for managing a storage environment comprising a plurality of storage units, the system comprising:

means for detecting a condition associated with a first storage unit from the plurality of storage units;

means for determining a first group from a plurality of groups to which the first storage unit belongs, wherein each group comprises one or more storage units from the plurality of storage units and inclusion of a storage unit in a group depends on a cost of storing data on the storage unit;

means for identifying a second group from the plurality of groups having an associated data storage cost that is lower than a data storage cost associated with the first group;

means for identifying a file stored on the first storage unit to be moved;

means for identifying a storage unit from the second group for storing the file; and

means for moving the identified file from the first storage unit to the storage unit from the second group that has been identified for storing the file.

53. A system for managing a storage environment comprising a plurality of storage units, the system comprising:

means for detecting a condition associated with a first storage unit from the plurality of storage units;

means for identifying a file stored on the first storage unit to be moved;

means for identifying a storage unit from the plurality of storage units for storing the identified file, wherein the data storage cost associated with identified storage unit is lower than a data storage cost associated with the first storage unit; and

means for moving the identified file from the first storage unit to the storage unit from the second group that has been identified for storing the file.